Quality Diversity

Generating a repertoire of good solutions

Jørgen Nordmoen, September 30, 2016

UNIVERSITY OF OSLO

Introduction

Why do we need "Quality Diversity"?

The basics of Evolutionary Algorithms (EAs)

An EA consist of:

- Population of solutions
- Fitness function
- Mutation and recombination operators

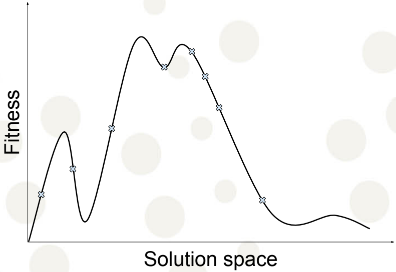

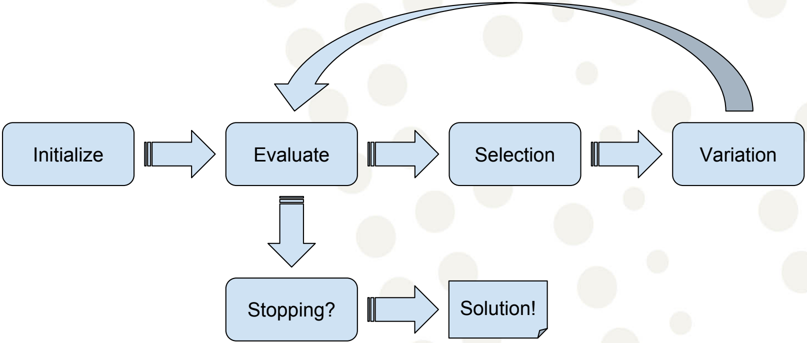
Initialize
Evaluate
Selection
Variation
Stopping?
Solution!

A solution in the EA:

- Represented with a *Genotype*
- Expressed with a *Phenotype*

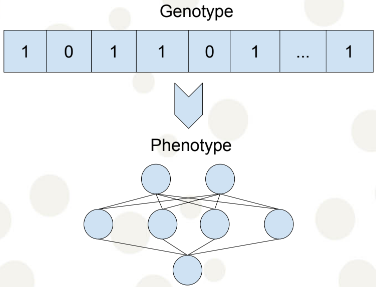

We apply the fitness function to the phenotype and the mutation and recombination operators to the genotype.

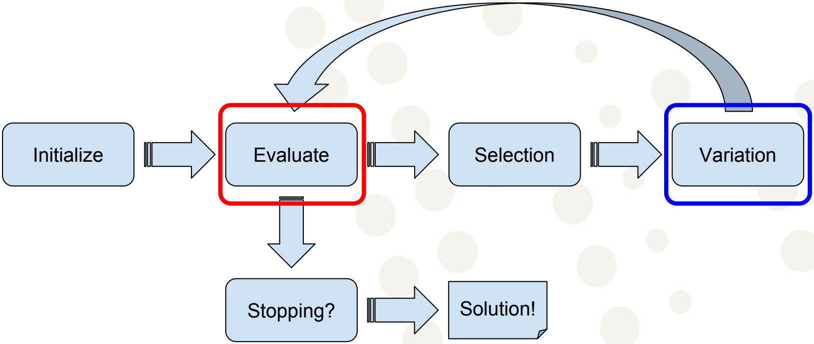
Initialize
Evaluate
Selection
Variation
Stopping?
Solution!

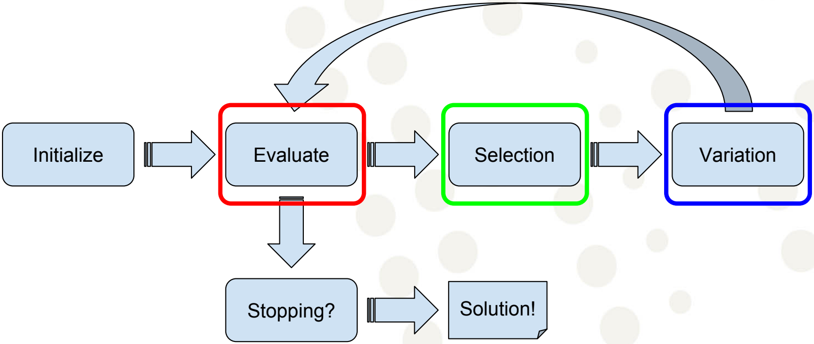
Initialize
Evaluate
Selection
Variation
Stopping?
Solution!

Some observations

- Convergent search (global optimization)
  - Selecting the "best" solutions
  - Homogenizing the "gene pool"

Some observations

- Convergent search (global optimization)
  - Selecting the "best" solutions
  - Homogenizing the "gene pool"
- Fitness and variation operates on different* representations
  - Evaluate phenotype
  - Change genotype

Some observations

- Convergent search (global optimization)
  - Selecting the "best" solutions
  - Homogenizing the "gene pool"
- Fitness and variation operates on different* representations
  - Evaluate phenotype
  - Change genotype
- Can't separate solutions with the same fitness
  - Selection does not differentiate how the fitness is expressed

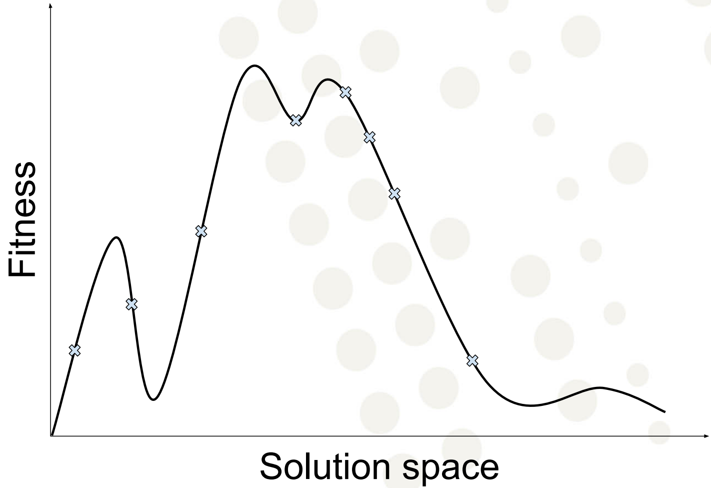
Fitness
Solution space

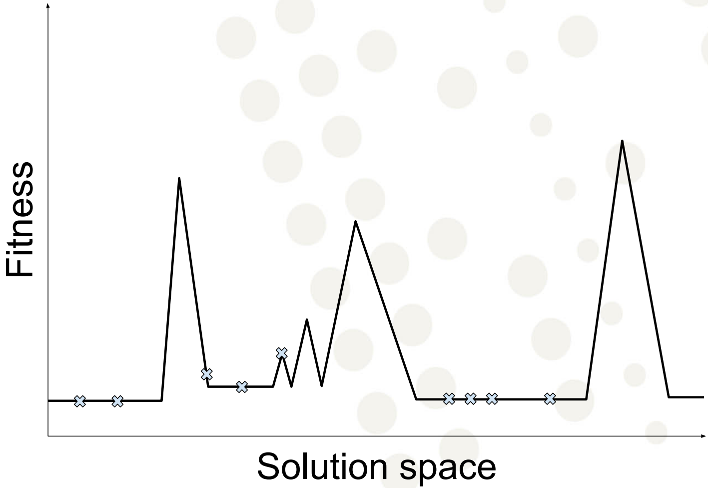
Fitness
Solution space


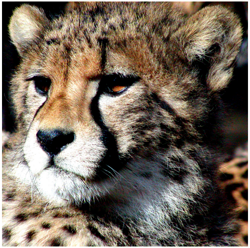


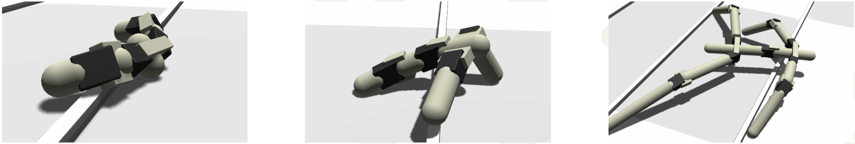

From Samuelsen and Glette[1]

Background

What is "Quality Diversity"?

Novelty Search[2]

- Introduction of divergent search
- Deceptive problems are hard for fitness based EA
- Searching for "interesting" solutions better than "good" solutions
- Archive of previous solutions

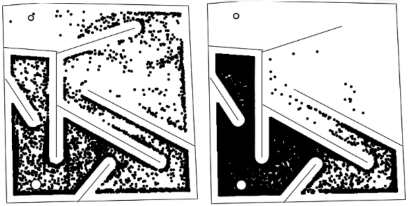

Behavioral diversity[3]

- Use expressed behavior as objective
  - Combine global fitness and behavior in multi-objective optimization
- Overcomes problems with similar fitness
- Expressed behavior can be simple and easy to calculate

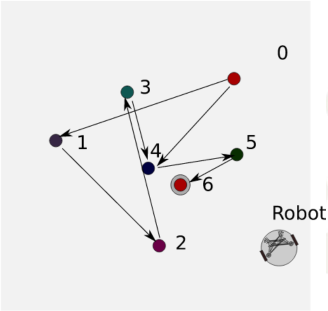

$$D_{ij} = ||v^{(i)} - v^{(j)}||$$

Novelty and behavioral diversity bears a striking resemblance

Novelty Search with Local Competition[4]

- Diversity on its own is not enough
- Global quality is still convergent
- Let niches evolve and have competition only inside niches!
  - Local competition allows each niche to become proficient
  - Novelty measured against whole population and archive

Novelty Search with Local Competition[4]

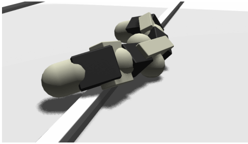

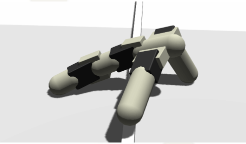

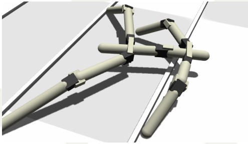

Behavioral Repertoire Learning in Robotics[5]

- Treat the novelty archive as a repertoire
- New solutions that are not novel, but better replaces older solutions
- Select desired behavior from repertoire as needed

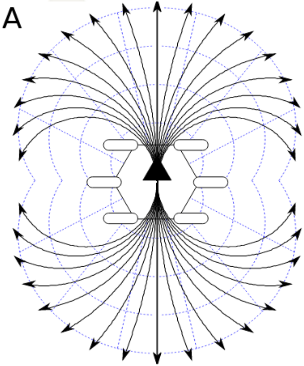

Illuminating search spaces by mapping elites[6]

- Divide behaviors into bins
- Behavior characteristics is the axis
- Fitness decides if a solution is put in a bin
- Easily search grid for best solution to current problem

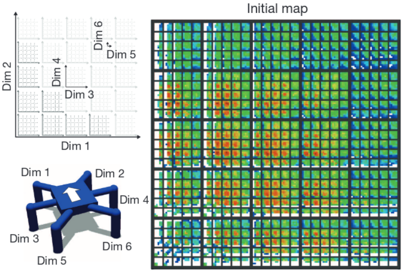

Summary

- Divergent search
- Quality among "similar" expressions
- "Behavior" is the important metric

"... the goal of this new type of search, called *quality diversity* (QD), is to find a maximally diverse collection of individuals (..) in which each member is as high performing as possible."[7]

Current research

Current research and challenges with Quality Diversity

Examples of real-world experiments

- Robots that can adapt like animals[8]
- Evolving a Behavioral Repertoire for a Walking Robot[9]
- EvoRBC: Evolutionary Repertoire-based Control for Robots with Arbitrary Locomotion Complexity[10]

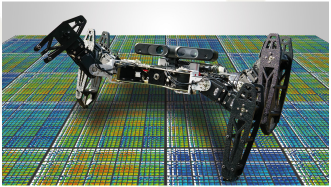

Examples of real-world experiments

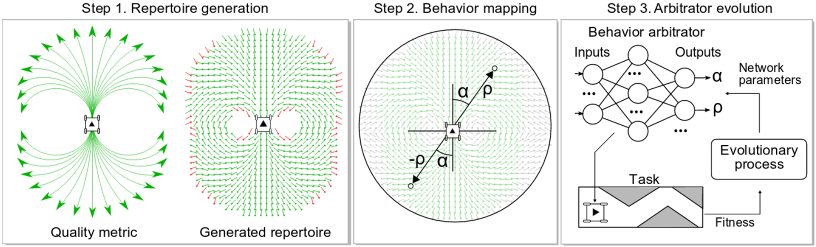

Encoding matters[11]

- Direct better than generative encoding (?)
- Generative encoding usually changes significantly
- Small influence on real-world tests!

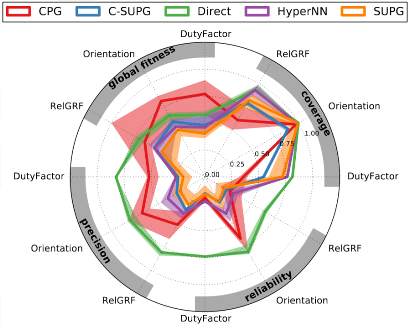

Alignment of behavioral characteristics (BC)

- Alignment is how well the BC is matched with fitness
- Pugh et al.[7] showed that
  - Choice of alignment is important in deceptive problems
  - Using multiple alignments can be used, but might not generate better results
- Auerbach et al.[12] showed that
  - BC drives the search in different directions
  - Novelty Search and MAP-Elites needs different types of alignment

Evolving generalists[13]

- Quality Diversity algorithms creates specialists
- Need to encourage each solution to acquire new skills

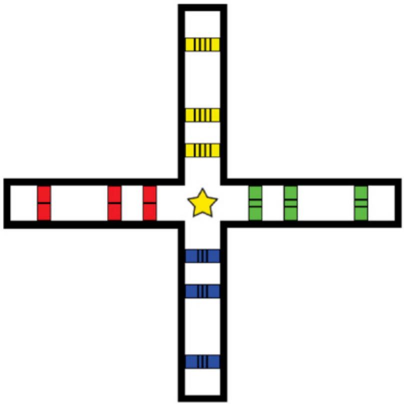

Open questions

- Isn't this just brute-force?
- How is this different from multi-objective optimization?
- Isn't this just moving all knowledge into behavior characteristics?
- Is Quality Diversity just a result that we can throw more compute power at the problem?
  - Can't we solve the same problem with regular EAs with larger populations?
  - Aren't the problems just so artificial that "Novelty" must win?

"Take home message"

- Quality Diversity is new way to think about searching
  - Diversity can create stepping stones to future solutions
  - Quality is important to direct the search

Bibliography

[1] Eivind Samuelsen and Kyrre Glette. Some distance measures for morphological diversification in generative evolutionary robotics. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*, pages 721–728. ACM, 2014.

[2] Joel Lehman and Kenneth O. Stanley. Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE*, pages 329–336, 2008.

[3] Jean-Baptiste Mouret and Stéphane Doncieux. Overcoming the bootstrap problem in evolutionary robotics using behavioral diversity. In *2009 IEEE Congress on Evolutionary Computation*, pages 1161–1168. IEEE, 2009.

[4] Joel Lehman and Kenneth O. Stanley. Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 211–218. ACM, 2011.

[5] Antoine Cully and Jean-Baptiste Mouret. Behavioral repertoire learning in robotics. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, pages 175–182. ACM, 2013.

[6] Jean-Baptiste Mouret and Jeff Clune. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*, 2015.

[7] Justin K. Pugh, Lisa B. Soros, and Kenneth O. Stanley. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI*, 3:40, 2016.

[8] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. Robots that can adapt like animals. *Nature*, 521(7553):503–507, 2015.

[9] Antoine Cully and J-B Mouret. Evolving a behavioral repertoire for a walking robot. *Evolutionary computation*, 24(1):59–88, 2016.

[10] Miguel Duarte, Jorge Gomes, Sancho Moura Oliveira, and Anders Lyhne Christensen. Evorbc: Evolutionary repertoire-based control for robots with arbitrary locomotion complexity. In *Proceedings of the 18th annual conference on Genetic and evolutionary computation*. ACM, 2016.

[11] Danesh Tarapore, Jeff Clune, Antoine Cully, and Jean-Baptiste Mouret. How do different encodings influence the performance of the MAP-Elites algorithm? In *Genetic and Evolutionary Computation Conference*, 2016.

[12] Joshua E. Auerbach, Giovanni Iacca, and Dario Floreano. Gaining insight into quality diversity. In *Proceedings of the 2016 on Genetic and Evolutionary Computation Conference Companion*, pages 1061–1064. ACM, 2016.

[13] Christopher Stanton and Jeff Clune. Curiosity search: Producing generalists by encouraging individuals to continually explore and acquire skills throughout their lifetime. *PloS one*, 11(9):e0162235, 2016.